

A Linear Approximation to the χ^2 Kernel with Geometric Convergence

Fuxin Li, *Member, IEEE*, Guy Lebanon, *Member, IEEE*, and Cristian Sminchisescu, *Member, IEEE*

Abstract

We propose a new analytical approximation to the χ^2 kernel that converges geometrically. The analytical approximation is derived with elementary methods and adapts to the input distribution for optimal convergence rate. Experiments show the new approximation leads to improved performance in image classification and semantic segmentation tasks using a random Fourier feature approximation of the $\exp -\chi^2$ kernel. Besides, out-of-core principal component analysis (PCA) methods are introduced to reduce the dimensionality of the approximation and achieve better performance at the expense of only an additional constant factor to the time complexity. Moreover, when PCA is performed jointly on the training and unlabeled testing data, further performance improvements can be obtained. Experiments conducted on the PASCAL VOC 2010 segmentation and the ImageNet ILSVRC 2010 datasets show statistically significant improvements over alternative approximation methods.

Index Terms



-
- Fuxin Li is with the School of Interactive Computing, Georgia Institute of Technology, Atlanta, GA, 30332.
E-mail: fli@cc.gatech.edu
 - Guy Lebanon is with the Department of Computational Science and Engineering, Georgia Institute of Technology, Atlanta, GA, 30332.
 - Cristian Sminchisescu is with the Center for Mathematical Sciences, Lund University, Lund, Sweden.

A Linear Approximation to the χ^2 Kernel with Geometric Convergence

1 INTRODUCTION

Histograms are important tools for constructing visual object descriptors. Many visual recognition approaches utilize similarity comparisons between histogram descriptors extracted from training and testing images. Widely used approaches such as k -nearest neighbors and support vector machines compare the testing descriptor with multiple training descriptors, and make predictions by a weighted sum of these comparison scores.

An important metric to compare histograms is the exponential- χ^2 kernel (referred to as $\exp -\chi^2$ in the rest of the paper), derived from the classic Pearson χ^2 test and utilized in many state-of-the-art object recognition studies [1], [2], [3], [4] with excellent performance. However, in the current big data era, training sets often contains millions to billions of examples. Training and testing via hundreds of thousands of comparisons using a nonlinear metric is often very time-consuming.

There are two main approaches to approximate the $\exp -\chi^2$ to facilitate fast linear time training and testing. One approach is to devise a transformation so that the χ^2 function can be represented as an inner product between two vectors. On top of this transformation, the random Fourier (RF) features methodology [5] is used to approximate a Gaussian kernel. The full $\exp -\chi^2$ kernel can be approximated by inner products on the vector after these two transformations [6]. A different approach is the Nyström method [7], which directly takes a subset of training examples, apply the comparison metric between an example and this subset and use the output as the feature vector (sometimes followed by principal component analysis (PCA)).

In this paper, we pursue the RF research line. We are interested in RF because it has the potential of representing more complicated functions than the Nyström approach, which is confined to summations of kernel comparisons and hard to approximate functions not of that type. However, RF has not been able to outperform Nyström so far, especially on image data with the $\exp -\chi^2$ approximation.

We believe that one reason for the suboptimal previous performance of RF in the $\exp -\chi^2$ kernel is the inaccuracy in the approximation of the χ^2 metric. A significant contribution of this paper is a new analytic series to approximate the χ^2 kernel. The new series is

derived using only elementary techniques and enjoys geometric convergence rate. Therefore, it is orders of magnitudes better in terms of approximation error than previously proposed approaches [8], [9]. Experiments show that this better approximation quality directly translate to better classification accuracy by using it in conjunction with the RF method to approximate the $\exp -\chi^2$ kernel.

Another research question we pursue is whether we can also improve the empirical performance of RF by applying PCA on the generated features. By applying PCA, the theoretical convergence rate of RF is no longer confined by the Monte Carlo rate $O(1/\sqrt{d})$, where d is the number of dimensions used in the approximation. Rather, it becomes dependent on eigenvalues, and with a fast enough eigenvalue decay rate, the convergence rate can reach $O(1/d)$ or better [10], raising it to the at least the same level as the Nyström approach. The question is then whether applying PCA on RF would translate to a comparable (or better) empirical performance.

For this question, we exploit out-of-core versions of PCA that add little computational overhead to the RF approximation, especially when combined with least squares and other quadratic losses, e.g. group LASSO. PCA allows us to reduce the number of dimensions required for classification and relaxes memory constraints when multiple kernels have to be approximated by RF. We also explore the use of unlabeled (test) data in order to better estimate the covariance matrix in PCA. This turns out to improve the performance by better selecting effective frequency components.

The paper is organized as follows: Section 2 summarizes related work. Section 3 describes the χ^2 kernel, where we elaborate the connection between the $\exp -\chi^2$ kernel and the χ^2 test. In Section 4, we present the new analytical approximation with geometric convergence rate. Section 5 describes the out-of-core PCA, Section 6 presents experiment results on PASCAL VOC 2010 and ImageNet ILSVRC 2010 data, and Section 7 concludes the paper.

2 RELATED WORK

To our knowledge, the use of the χ^2 kernel for histogram comparison can be traced back to at least 1996 [11]. [12] constructed the $\exp -\chi^2$ kernel and used it in SVM-based image classification. They hypothesized that exponential χ^2 is a Mercer kernel,

but the real proof was not available until 2004 in the appendix of [13]. The χ^2 kernel and $\exp -\chi^2$ has been used in a number of visual classification and object detection systems [1], [2], [3], [4] and has been shown to have one of the best performances among histogram kernels [14]. [15] proposes an extension to the χ^2 kernel that normalizes the χ^2 cross different bins. Other metrics for histogram comparison include histogram intersection, where an efficient speed-up for testing has also been proposed [16], Hellinger kernel, earth mover distance [17] and Jensen-Shannon. See [8] for a summary and comparisons.

Random Fourier features were proposed by [5] on translation-invariant kernels. [6] generalizes it to the $\exp -\chi^2$ kernel by the aforementioned two-steps approach. Several other studies on linear kernel approximations also used ideas in RF [8], [9], [18].

The Nyström method [7] sub-samples the training set and operate on a reduced kernel matrix. Its asymptotic convergence rate had long known to be slow [19], but recent papers have proved that it is actually faster than the Monte Carlo rate of RF [20]. Other speed-ups to kernel methods based on low-rank approximations of the kernel matrix have been proposed in [21], [22].

A topic of recent interest is methods for coding image features, where the goal is to achieve good performance using linear learners following a feature embedding [23], [24]. Hierarchical coding schemes based on deep structures have also been proposed [25]. Both sparse and dense coding schemes have proven successful, with supervector coding [26] and the Fisher kernels [27] some of the best performers in the ImageNet large-scale image classification challenge [28]. The dictionaries of some influential coding schemes are usually extremely large – both the Fisher kernel and supervector coding usually require more than 200k dimensions [29]) and the training of the dictionary is often time-consuming. RF and Nyström do not require training, hence they are interesting alternatives to these methods.

A crucial component in many coding algorithms is a max-pooling approach, which uses the maximum of the coded descriptors in a spatial range as features. Since in this case an informative small patch could have the same descriptor as the whole image, it is desirable in image classification (for highlighting important regions) but undesirable for object detection and semantic segmentation problems, where the size and shape of the object is of interest. A recent second-order pooling scheme [30] proposes an alternative and has shown successful results in the semantic segmentation problem.

3 THE χ^2 KERNEL AND ITS RELATIONSHIP WITH THE χ^2 TEST

◦ denotes element-wise products of vectors. $\frac{a}{b}$ denotes an element-wise division of b from a .

The χ^2 kernel is derived from Pearson's χ^2 test. The original Pearson χ^2 test is for testing whether an empirical histogram estimate matches a probability distribution. Given a histogram estimate $\mathbf{x} = [x_1, x_2, \dots, x_d]$, the test statistic is

$$X^2(\mathbf{x}, \mathbf{E}) = \sum_{i=1}^d \frac{(x_i - E_i)^2}{E_i} \quad (1)$$

where $\mathbf{E} = [E_1, E_2, \dots, E_d]$ is the theoretical frequency in the bins.

Suppose we have two histogram estimates \mathbf{x} and \mathbf{y} , one can arrive at a symmetric version by taking the harmonic mean $H(x, y) = \frac{2}{1/x + 1/y}$ of each bin and sum it up:

$$\begin{aligned} \chi^2(x, y) &= \frac{1}{4} \sum_{i=1}^d H\left(\frac{(x_i - y_i)^2}{y_i}, \frac{(y_i - x_i)^2}{x_i}\right) \\ &= \frac{1}{2} \sum_{i=1}^d \frac{(x_i - y_i)^2}{x_i + y_i} \end{aligned} \quad (2)$$

The virtue of such a harmonic mean approach lies in removing the singular points in the kernel: the value of the original χ^2 test goes to infinity when $E_i = 0$ and $x_i \neq 0$. Using the harmonic mean approach in (2), the function is well-defined in all cases.

In order to use the X^2 test (1) to determine goodness of the fit, one needs to compute the p-value of the X^2 statistic:

$$p = 1 - P\left(\frac{k}{2}, \frac{X^2}{2}\right) \quad (3)$$

where k is the degree of freedom in the distribution, $P(k, x)$ is the regularized Gamma function. The p-value is 1 minus the cumulative distribution function (CDF) the test statistic. If a p-value is small, then it means the observed statistic is very unlikely to happen under the hypothesized distribution. A usual criterion is to decide that \mathbf{x} disagrees from the distribution specified by $[E_1, \dots, E_d]$ if $p < 0.05$. In the case of the χ^2 test, with a special case of $k = 2$, one has $p = \exp(-\frac{X^2}{2})$, which coincides with the $\exp -\chi^2$ kernel. Since the p-value is the relevant metric for comparing two distributions, the $\exp -\chi^2$ kernel can be considered intuitively better than the χ^2 function as a similarity metric comparing two histogram distributions. Empirically, we have tested kernels with different degrees of freedom, and found out that $\exp -\chi^2$ works similarly to $\text{erf}(\chi^2)$ (corresponding to χ^2 with 1 degree of freedom) while outperforming all others with more than 2 degrees of freedom.

4 AN ANALYTICAL APPROXIMATION TO THE χ^2 KERNEL WITH GEOMETRIC CONVERGENCE

In the following we show an analytical approximation to the χ^2 kernel. We start with the one-dimensional

case. The χ^2 kernel in one dimension has the form:

$$\chi^2(x_i, y_i) = \frac{(x_i - y_i)^2}{2(x_i + y_i)} = \frac{1}{2}(x_i + y_i) - \frac{2x_i y_i}{x_i + y_i} \quad (4)$$

Because $\sum_i x_i = 1$ and $\sum_i y_i = 1$ in a histogram, the first $x_i + y_i$ form sums to a constant. It is thus important to represent the form $\frac{xy}{x+y}$ into an inner product. We will make repeated use of the following crucial formula

$$\begin{aligned} \frac{2xy}{x+y} &= \frac{x-k}{x+k} \frac{y-k}{y+k} \frac{2xy}{x+y} + \left(1 - \frac{x-k}{x+k} \frac{y-k}{y+k}\right) \frac{2xy}{x+y} \\ &= \frac{x-k}{x+k} \frac{y-k}{y+k} \frac{2xy}{x+y} + \frac{2\sqrt{kx}}{x+k} \frac{2\sqrt{ky}}{y+k} \end{aligned} \quad (5)$$

which gives a one-term linear approximation of the χ^2 kernel with $c_x = \frac{2\sqrt{kx}}{x+k}$. Repeatedly plugging (5) into the $\frac{2xy}{x+y}$ in the first term of the right-hand-side in (5) gives us a series with multiple parameters:

$$c_x = \left[\frac{2\sqrt{k_1x}}{x+k_1}, \frac{x-k_1}{x+k_1} \frac{2\sqrt{k_2x}}{x+k_2}, \frac{x-k_1}{x+k_1} \frac{x-k_2}{x+k_2} \frac{2\sqrt{k_3x}}{x+k_3}, \dots \right] \quad (6)$$

This series has geometric convergence rate as the N -term error is exactly:

$$\frac{(x-k_1)\dots(x-k_N)(y-k_1)\dots(y-k_N)}{(x+k_1)\dots(x+k_N)(y+k_1)\dots(y+k_N)} \frac{2xy}{x+y} \quad (7)$$

which is straightforwardly geometric if we take $k = k_1 = \dots = k_N$, because $|\frac{x-k}{x+k}| < 1, \forall 0 < k \leq 1$.

We see the multiple parameters k_1, k_2, \dots in this series a boon rather than a distraction, because it can greatly improve the convergence rate in the full range of $[0, 1]$. Note that the convergence rate is dominated by $\left(\frac{x-k}{x+k}\right)^{2N}$ if there is only one $k = k_1 = \dots = k_N$. Although this is fast in general, it can be very slow if $\frac{x-k}{x+k}$ is close to 1. Two examples are: $k = 1, x = 0.005$ and $k = 0.005, x = 1$. Apparently, there is no single k choice that achieves good convergence rate on the entire input domain $[0, 1]$. Our solution is to utilize multiple different parameters to cover different regions, and combining the parameter choice with the input distribution of our data to achieve an optimal convergence rate on the entire domain of the input.

First we establish a simple upper bound of the function to facilitate simpler error computation:

$$\frac{2xy}{x+y} \leq \frac{2x}{x+1} \quad (8)$$

Now the N -term error can be represented as:

$$\frac{2xy}{x+y} - c_x^\top c_y \leq \frac{2(x-k_1)\dots(x-k_N)x}{(x+k_1)\dots(x+k_N)(x+1)} \quad (9)$$

Our algorithm for finding the parameters proceeds greedily to eliminate the highest error peak at each iteration. Specifically, we choose the parameter:

$$k_{N+1} = \arg \max_x \left| \frac{2(x-k_1)\dots(x-k_N)x}{(x+k_1)\dots(x+k_N)(x+1)} p(x) \right| \quad (10)$$

where $p(x)$ is the input distribution of x , estimated on each particular dataset. Such a choice reduces error to 0 at the mode of the input distribution and is empirically tested to be superior than other greedy schemes such as minimizing the mean error at each step. In practice, $p(x)$ is estimated using a histogram estimate with logarithmically spaced bins, and k_{N+1} is chosen as one of the bin centers. The algorithm of such an implementation is shown in Algorithm 1.

Note that the χ^2 kernel in this form coincides with the harmonic mean of the two vectors. Therefore, our approach could also be a linear approximation on the harmonic mean between two vectors. However, currently we do not know of applications of that.

Algorithm 1 Find the parameters for the input distribution specified by feature matrix \mathbf{X} .

input : Feature matrix \mathbf{X} , Parameter vector length N .

output : parameter vector \mathbf{k} .

- 1: Compute a histogram density estimate \mathbf{h} on all nonzero values in \mathbf{X} using logarithmically spaced bins in the range $[\min_{x \in \mathbf{X}, x \neq 0} \mathbf{X}, \max \mathbf{X}]$, denote the vector of bin centroids as \mathbf{x} .
 - 2: $\mathbf{b} = \frac{\mathbf{x}}{\mathbf{x}+1} \circ \mathbf{h}$
 - 3: **for** $i = 1 \rightarrow N$ **do**
 - 4: $k_i = x_j, j = \arg \max_j |b_j|$
 - 5: $\mathbf{b} = \mathbf{b} \circ \frac{\mathbf{x}-k_i}{\mathbf{x}+k_i}$
 - 6: **end for**
-

5 PRINCIPAL COMPONENT ANALYSIS OF RANDOM FEATURES ON MULTIPLE DESCRIPTORS

Another rather orthogonal strategy we pursue is principal component analysis after obtaining random features, and solving regression problems after the PCA. Care needs to be exercised when PCA is performed on an extremely large-scale dataset in conjunction with multiple kernels. Although similar approaches have been discussed extensively in the high-performance computing literature ((e.g., [31]), we have not found such treatments detailed in the vision community, especially in the context of linear approximations for kernel methods.

The main advantage of using PCA after RF (hereafter called RF-PCA) is to reduce the memory footprint. It is known that the performance of RF improves when more random dimensions are used. However, the speed of learning algorithms usually deteriorates quickly when the data cannot be load in memory, which would be the case when the RF of multiple kernels are concatenated: e.g. with 7 kernels and 7,000 RF dimensions for each kernel, the learning phase following RF needs to operate on a 49,000 dimensional feature vector.

Using eigenvectors is also one of the very few approaches that could provide a better asymptotic

convergence rate than the $O(\frac{1}{\sqrt{m}})$ for Monte Carlo, thus requiring fewer dimensions for an approximation of the same quality. Many other techniques like quasi-Monte Carlo suffer from the curse of dimensionality – the convergence rate decreases exponentially with the number of input dimensions [32], which generally makes it unsuitable for RF which is supposed to work on high-dimensional problems.

Another interesting aspect of RF-PCA is it can bring an unexpected flavor of semi-supervised learning, in that one can use unlabeled test data to improve classification accuracy. RF-PCA amounts to selecting the relevant dimensions in the frequency domain, by considering both the training and testing data during PCA, frequencies that help discriminate test data will more likely be selected. In the experiments such a strategy will be shown to improve performance over the computation of PCA only on training data.

One main problem is, in a large training set, the feature matrix cannot be fully loaded into memory. Therefore PCA needs to be performed out-of-core, a high-performance computing term depicting this situation (unable to load data into memory). The way to do PCA in linear time is not by singular value decomposition on the RF features Z , but rather by performing eigenvalue decomposition for the centered covariance matrix $Z^T(I - \frac{1}{n}\mathbf{1}\mathbf{1}^T)Z$. $Z^T Z = \sum_i Z_i^T Z_i$ can be computed out-of-core by just loading a chunk of X_i into memory at a time, compute their RF feature Z , compute the covariance matrix and then delete the RF features from memory. Then an eigen-decomposition gives the transformation matrix U for PCA. We denote \bar{U} as the matrix obtained by selecting the first D dimensions of U corresponding to the largest eigenvalues. Denote the mean vector of the input matrix $\bar{Z} = \frac{1}{n}Z^T \mathbf{1}$, and

$$\tilde{Z} = (Z - \mathbf{1}\bar{Z}^T)\bar{U} = (I - \frac{1}{n}\mathbf{1}\mathbf{1}^T)Z\bar{U} \quad (11)$$

is the feature vector obtained after PCA projection (Algorithm 2). It is very convenient to perform regres-

sion with a quadratic loss after PCA, since only the Hessian is needed for optimization. This applies not only to traditional least squares regression, but also to the LASSO, group LASSO, and other composite regularization approaches. In this case the projections need not be performed explicitly. Instead, notice that only $\tilde{Z}^T \tilde{Z}$ and $\tilde{Z}^T y$ are needed for regression:

$$\begin{aligned} \tilde{Z}^T \tilde{Z} &= \bar{U}^T Z^T (I - \frac{1}{n}\mathbf{1}\mathbf{1}^T) Z \bar{U} \\ \tilde{Z}^T y &= \bar{U}^T Z^T (I - \frac{1}{n}\mathbf{1}\mathbf{1}^T) y \end{aligned} \quad (12)$$

It follows that only $Z^T Z$, $Z^T \mathbf{1}$ and $Z^T y$ have to be computed. All terms can be computed out-of-core simultaneously. Algorithm 3 depicts this scenario.

Algorithm 3 Learning after PCA with Quadratic Loss.

input : $n \times d$ data matrix $X = [X_1^T, X_2^T, \dots, X_n^T]^T$.

Output vector y . Number of dimension D to retain after PCA.

- 1: Perform out-of-core PCA using Algorithm 2.
 - 2: $H' = \bar{U}^T H \bar{U} = \bar{D}$, the first D rows and columns of the diagonal matrix \mathbf{D} .
 - 3: $v' = \bar{U}^T v - \frac{1}{n}(\mathbf{1}^T y)\bar{U}^T m$.
 - 4: Perform learning on H', v' , e.g., for linear ridge regression where the optimization is $\arg \min_w \|w^T \tilde{Z} - y\|^2 + \lambda \|w\|^2$, the solution is $w = (H' + \lambda I)^{-1} v'$.
 - 5: Use $\bar{U}^T w$ instead of w as a function of the original inputs: $f(x) = w^T \bar{U} x - \frac{1}{n} w^T \bar{U} m$, in order to avoid the projection for the testing examples.
-

Under this PCA approach the data is loaded only once to compute the Hessian. Additional complexity of $O(D^3)$ is necessary for matrix decomposition on H . If ridge regression is used, the H' after decomposition is diagonal therefore only $O(D)$ is needed to obtain the regression results. In this case the additional constant factor is quite small. The bottleneck of this algorithm for large-scale problems is undoubtedly the computation of the initial Hessian, which involves reading multiple chunks from disk.

The more sophisticated case is when PCA needs to be performed separately on multiple different kernel approximators, i.e., $Z = [Z^{(1)} Z^{(2)} \dots Z^{(l)}]$, where each $Z^{(i)}$ is the RF feature embedding of each kernel. This time, the need to compute $Z^{(i)T} Z^{(j)}$ rules out tricks for simple computation. The data needs to be read in twice (Algorithm 4), first to perform the PCA, and then use U to transform X in chunks in order to obtain Z and $Z^T Z$. But the full computation is still linear in the number of training examples. In both cases, the projection is not required for the testing examples. Because whenever w is obtained, $w^T \tilde{Z} = w^T \bar{U} (Z - \frac{1}{n} Z \mathbf{1}^T)$, then $\bar{U} w$ can be the weight vector for the original input, with the addition of a constant term.

Algorithm 2 Out-of-Core Principal Component Analysis.

input : $n \times d$ data matrix $X = [X_1^T, X_2^T, \dots, X_n^T]^T$.

Output vector y . Number of dimension D to retain after PCA.

- 1: Divide the data into k chunks, called $X_{(1)}, X_{(2)}, \dots, X_{(k)}$.
 - 2: $H = \mathbf{0}, m = \mathbf{0}, v = \mathbf{0}$
 - 3: **for** $i = 1 \rightarrow k$ **do**
 - 4: Load the i -th chunk $X_{(i)}$ into memory.
 - 5: Use Algorithm ?? to compute the RF feature $Z_{(i)}$ for $X_{(i)}$.
 - 6: $H = H + Z_{(i)}^T Z_{(i)}, m = m + Z_{(i)}^T \mathbf{1}, v = v + Z_{(i)}^T y$
 - 7: **end for**
 - 8: $H = H - \frac{1}{n} m m^T$.
 - 9: Compute eigen-decomposition $H = U \mathbf{D} U^T$. Output the first D columns of U as \bar{U} , the diagonal matrix \mathbf{D} , and the input-output product v .
-

Algorithm 4 Two-stage Principal Component Analysis when learning with multiple kernels.

input : $n \times d$ data matrix $X = [X_1^T, X_2^T, \dots, X_n^T]^T$.
 Output vector y . Number of dimension D to retain after PCA.

- 1: Perform out-of-core PCA using Algorithm 2.
 - 2: **for** $i = 1 \rightarrow k$ **do**
 - 3: Load the i -th chunk $X_{(i)}$ into memory.
 - 4: Use Algorithm 1 to compute the RF feature $Z_{(i)}$ for $X_{(i)}$, with the same randomization vectors w as before.
 - 5: $\tilde{Z} = (Z_{(i)} - \frac{1}{n} \mathbf{1} m^T) \bar{U}$.
 - 6: $H' = H' + \tilde{Z}^T \tilde{Z}$, $v' = v' + \tilde{Z}^T y$
 - 7: **end for**
 - 8: Perform learning on H', v' . E.g., for linear least squares where the optimization is $\arg \min_w \|w^T Z - y\|^2$, the solution is $w = H'^{-1} v'$.
 - 9: Use $\bar{U}^T w$ instead of w as a function of the original inputs: $f(x) = w^T \bar{U} x - \frac{1}{n} w^T \bar{U} m$, in order to avoid the projection step for the testing examples.
-

It is worth noting that out-of-core least squares or ridge regression scales extremely well with the number of output dimensions c , which can be used to solve one-against-all classification problems with c classes. In the out-of-core case, $Z^T y$ will be computed in $O(nDc)$ time along with the Hessian in Algorithm 2 or 4. After the inverse of Hessian is obtained, only a matrix-vector multiplication costing $O(D^2c)$ is needed to obtain all the solutions, without any dependency on n . Thus the total time of this approach with c classes is $O(nDc + D^2c)$ which scales very nicely with c . Especially compared with other algorithms that need to perform the full training procedure on each class. Although the L_2 loss is not optimal for classification, in large-scale problems (e.g. ImageNet) with 1,000 – 10,000 classes, the out-of-core ridge regression can still be used to generate a fairly good baseline result quickly.

6 EXPERIMENTS

Our experiments are conducted on two challenging datasets: PASCAL VOC 2010 [33] and ImageNet [28] ILSVRC 2010 (<http://www.image-net.org/challenges/LSVRC/2010/>). These benchmarks reveal the performance differences among approximation methods, which would otherwise be difficult to observe in simple datasets. We conduct most experiments on the medium-scale PASCAL VOC data in order to compare against exact kernel methods. For this dataset, we use exclusively the `train` and `val` datasets, which have 964 images and around 2100 objects each. Classification results are also shown on the ImageNet dataset to demonstrate the efficiency of our kernel approximations. The experiments are conducted




Fig. 1: Comparisons on various approximations to the χ^2 kernel. It can be seen that the new direct approximation is converging orders of magnitude faster than previous approaches.

using an Intel Xeon E5520 2.27GHz with 8 cores and 24GB memory. The algorithm ?? is parallelized using OpenMP to take advantage of all cores.

6.1 Comparing Approximations

To test the different approximation, we consider a medium-scale problem from the PASCAL VOC segmentation dataset. For training, we use image segments (obtained using the constrained parametric min-cuts algorithm, CPMC [34]) that best match each ground truth segment in terms of overlap (called best-matching segments) in the `train` set, plus the ground truth segments. The best-matching segments in the `val` set are used as test. This creates a medium-scale problem with 5100 training and 964 test segments.

The methods tested in experiments are Chebyshev, VZ [8], Direct. For reference, we also report classification results for the χ^2 kernel without exponentiating as `Chi2`, as well as the skewed χ^2 kernel proposed in [18] as `Chi2-Skewed`. Due to the Monte Carlo approximation, different random seeds can lead to quite significant performance variations. Therefore the experiments are all averaged over 50 trials on random seeds. Within each trial, the same random seeds are used for all methods. For PCA-Chebyshev, the initial sampling is done using three times the final approximating dimensions, and PCA is performed to reduce the dimensionality to the same level as the other two methods. We test the classification performance of these kernels with two different types of features: a bag of SIFT words (BOW) feature of 300 dimensions, and a histogram of gradient (HOG) feature of 1700 dimensions. The classification is done via a linear SVM using the LIBSVM library (empirically we found the LIBLINEAR library produced worse results than LIBSVM in this context with dense features). The C parameter in LIBSVM is validated to 50, the kernel to be approximated is $\exp(-\chi^2)$, with $\beta = 1.5$. For VZ, the period parameter is set to the optimal one specified in [8]. For each kernel, 5 dimensions are used to approximate the χ^2 distance in each dimension, which represents a common use case.

Number of Dimensions	3000	5000	7000
Chi2	41.91%	42.32%	42.12%
Chi2-Skewed	39.82% \pm 0.73%	40.79% \pm 0.55%	40.90% \pm 0.82%
Chebyshev	42.03% \pm 0.80%	42.61% \pm 0.65%	42.61% \pm 0.64%
VZ	42.29% \pm 0.74%	42.82% \pm 0.63%	43.00% \pm 0.57 %
Direct	42.09% \pm 0.79%	42.88% \pm 0.63%	43.21% \pm 0.63%
PCA-Chebyshev	42.80% \pm 0.74%	43.25% \pm 0.55%	43.42% \pm 0.42 %
PCA-VZ			
PCA-Direct	43.16% \pm 0.55%	43.31% \pm 0.53%	43.53% \pm 0.71 %
Exact $\exp-\chi^2$	44.19%		

TABLE 1: Classification accuracy of $\exp-\chi^2$ kernel when the χ^2 function is estimated with different approximations, on a BOW-SIFT descriptor. Results for the Chi2 and Chi2-Skewed kernels are also shown for reference.

Number of Dimensions	3000	5000	7000
Chi2	29.15%	30.50%	31.22%
Chi2-Skewed	30.08% \pm 0.74%	30.37 % \pm 0.63%	30.51 % \pm 0.35 %
Chebyshev	30.86% \pm 0.78%	31.53% \pm 0.66%	31.90% \pm 0.70%
VZ	31.32% \pm 0.90%	32.07 % \pm 0.83%	32.36% \pm 0.62%
Direct	31.71% \pm 0.92%	32.72% \pm 0.73%	32.94% \pm 0.66%
PCA-Chebyshev	32.59% \pm 0.77%	33.11% \pm 0.57%	33.22% \pm 0.54%
PCA-VZ	32.94% \pm 0.67%	33.41% \pm 0.54%	33.45% \pm 0.59%
PCA-Direct	32.92% \pm 0.66%	33.68% \pm 0.57%	33.63% \pm 0.67%
Exact $\exp-\chi^2$	34.34%		

TABLE 2: Classification accuracy of $\exp-\chi^2$ kernel when the χ^2 function is approximated with different approximations, on a HOG descriptor. Results for the Chi2 and Chi2-Skewed kernels are also shown for a reference.




Fig. 2: Effect on the classification accuracy on a 7000-dimensional RF-approximated $\exp-\chi^2$ kernel, using different approximations and various number of dimensions to approximate the χ^2 function.

6.2 Results for Multiple Kernels on the PASCAL VOC Segmentation Challenge

In this section, we consider the semantic segmentation task from PASCAL VOC, where we need to both recognize objects in images, and generate pixel-wise segmentations for these objects. Ground truth segments of objects paired with their category labels are available for training.

A recent state-of-the-art approach trains a scoring function for each class on many putative figure-ground segmentation hypotheses, obtained using CPMC [?], [35]. This creates a large-scale learning task even if the original image database has moderate size: with 100 segments in each image, training for 964 images creates a learning problem with around 100,000 training examples. This training set is still tractable for exact kernel approaches and we can

directly compare against them.

Two experiments are conducted using multiple kernel approximations for the $\exp-\chi^2$ kernels. We use 7 different image descriptors, which include 3 HOGs at different scales, BOW on SIFT for the foreground and background, and BOW on color SIFT for the foreground and background [34], [36]. The VOC segmentation measure is used to compare the different approaches. This measure is the average of pixel-wise average precision on the 20 classes plus background. To avoid distraction and for a fair comparison, the post-processing step [34] is not performed and the result is obtained by only reporting one segment with the highest score in each image. The method used for nonlinear estimation is one-against-all support vector regression (SVR) as in [36], and the method for linear estimation is one-against-all ridge regression. The latter is used since fast solutions for linear SVR problems are not yet available for out-of-core dense features. We avoided stochastic gradient methods (e.g., [26]) since these are difficult to tune to convergence, and such effects can potentially bias the results. We average over 5 trials of different random seeds.

6.3 Results on ImageNet

The ImageNet ILSVRC 2010 is a challenging classification dataset where 1 million images have to be separated into 1,000 different categories. Here we only show experiments performed using the original BOW feature provided by the authors. Our goal is primarily to compare among different approximations, hence we did not generate multiple image descriptors or

Number of Dimensions	3000	5000	7000
Chebyshev	16.30% \pm 0.04%	17.11% \pm 0.04%	17.63% \pm 0.09%
PCA-Chebyshev	16.66% \pm 0.08%	17.85% \pm 0.08%	18.65% \pm 0.10 %
VZ	16.10% \pm 0.04%	16.95 % \pm 0.08%	17.48% \pm 0.09%
Direct			
Nyström			
Linear	11.6% ([37])		

TABLE 4: Performance of linear classifiers as well as non-linear approximation methods on ImageNet ILSVRC 2010 data. Notice the significant boost provided by the non-linear approximations (exact non-linear calculations are intractable in this large scale setting).

- [28] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2009.
- [29] K. Chatfield, V. Lempitsky, A. Vedaldi, and A. Zisserman, "The devil is in the details: an evaluation of recent feature encoding methods," in *Proceedings of the British Machine Vision Conference*, 2011.
- [30] J. Carreira, R. Caseiro, J. Batista, and C. Sminchisescu, "Semantic segmentation with second-order pooling," in *European Conference on Computer Vision*, 2012.
- [31] Y. Qu, G. Ostrouchov, N. Samatova, and A. Geist, "Principal component analysis for dimension reduction in massive distributed data sets," in *ICDM*, 2002.
- [32] R. Caflisch, "Monte carlo and quasi-monte carlo methods," *Acta Numerica*, vol. 7, pp. 1–49, 1998.
- [33] M. Everingham, L. V. Gool, C. Williams, J. Winn, and A. Zisserman, "The pascal visual object classes (voc) challenge," *International Journal of Computer Vision*, vol. 88, pp. 303–338, 2010.
- [34] J. Carreira and C. Sminchisescu, "CPMC: Automatic Object Segmentation Using Constrained Parametric Min-Cuts," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, July 2012.
- [35] —, "Constrained parametric min cuts for automatic object segmentation," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2010.
- [36] F. Li, J. Carreira, and C. Sminchisescu, "Object recognition as ranking holistic figure-ground hypotheses," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2010.
- [37] J. Deng, A. C. Berg, K. Li, and L. Fei-Fei, "What does classifying more than 10,000 image categories tell us?" in *European Conference on Computer Vision*, 2010.

[illegible]

Michael Shell Biography text here.

PLACE
PHOTO
HERE

John Doe Biography text here.Biography text here.Biography
text here.Biography text here.Biography text here.Biography
text here.Biography text here.Biography text here.Biography
text here.Biography text here.Biography text here.Biography
text here.Biography text here.Biography text here.Biography
text here.Biography text here.Biography text here.Biography
text here.Biography text here.Biography text here.Biography
text here.Biography text here.Biography text here.Biography
text here.Biography text here.Biography text here.
here.Biography text here.Biography text here.